

El Capitan: The NNSA Exascale Machine

Ian Karlin
Principal HPC Strategist Livermore Computing
El Cap COE Deputy

September 2nd, 2020




LLNL-PRES-787929

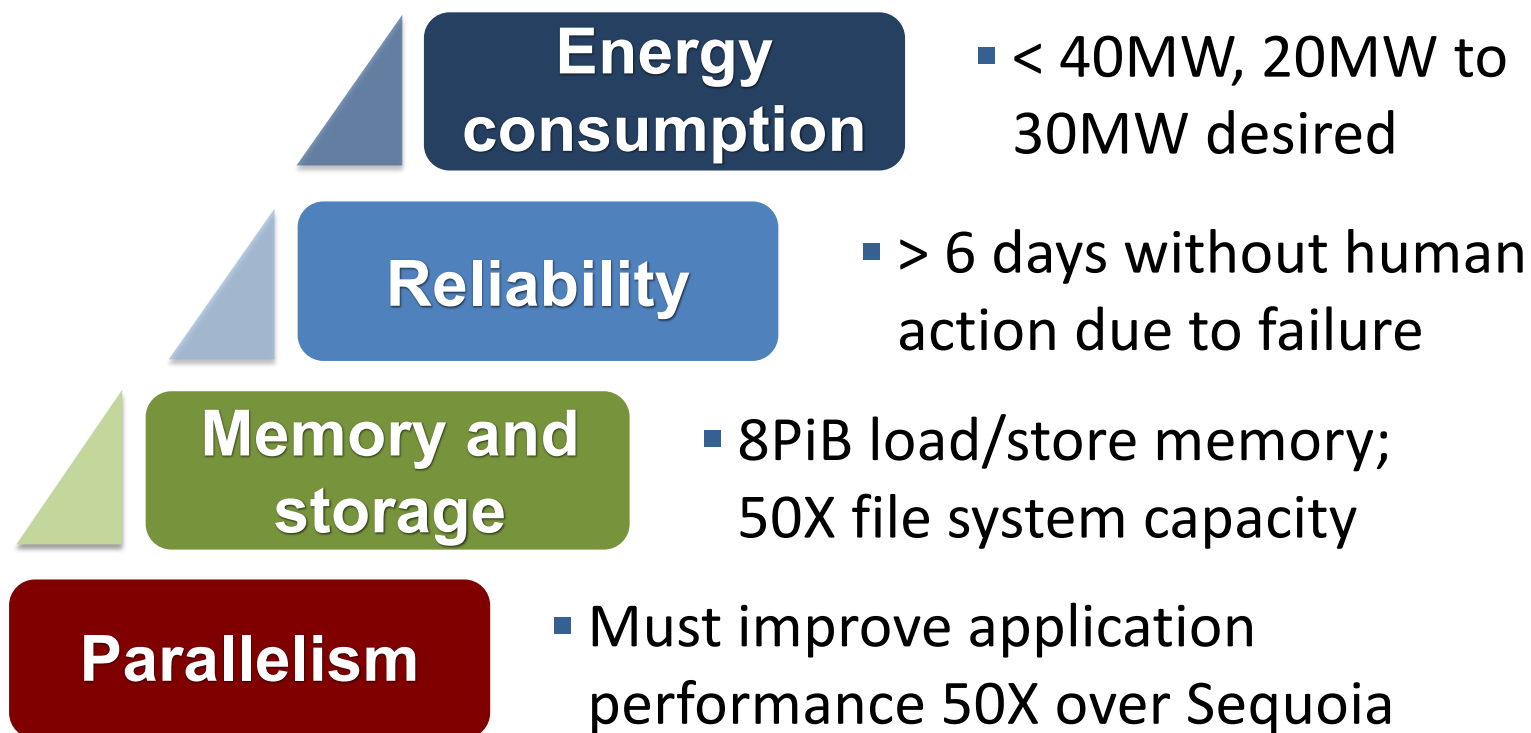
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

 Lawrence Livermore
National Laboratory

El Capitan will be LLNL's next ASC ATS platform



The CORAL2 RFP targets exascale systems



El Capitan can meet all of these high-level requirements

Cray and AMD will deliver a highly capable GPU-accelerated system

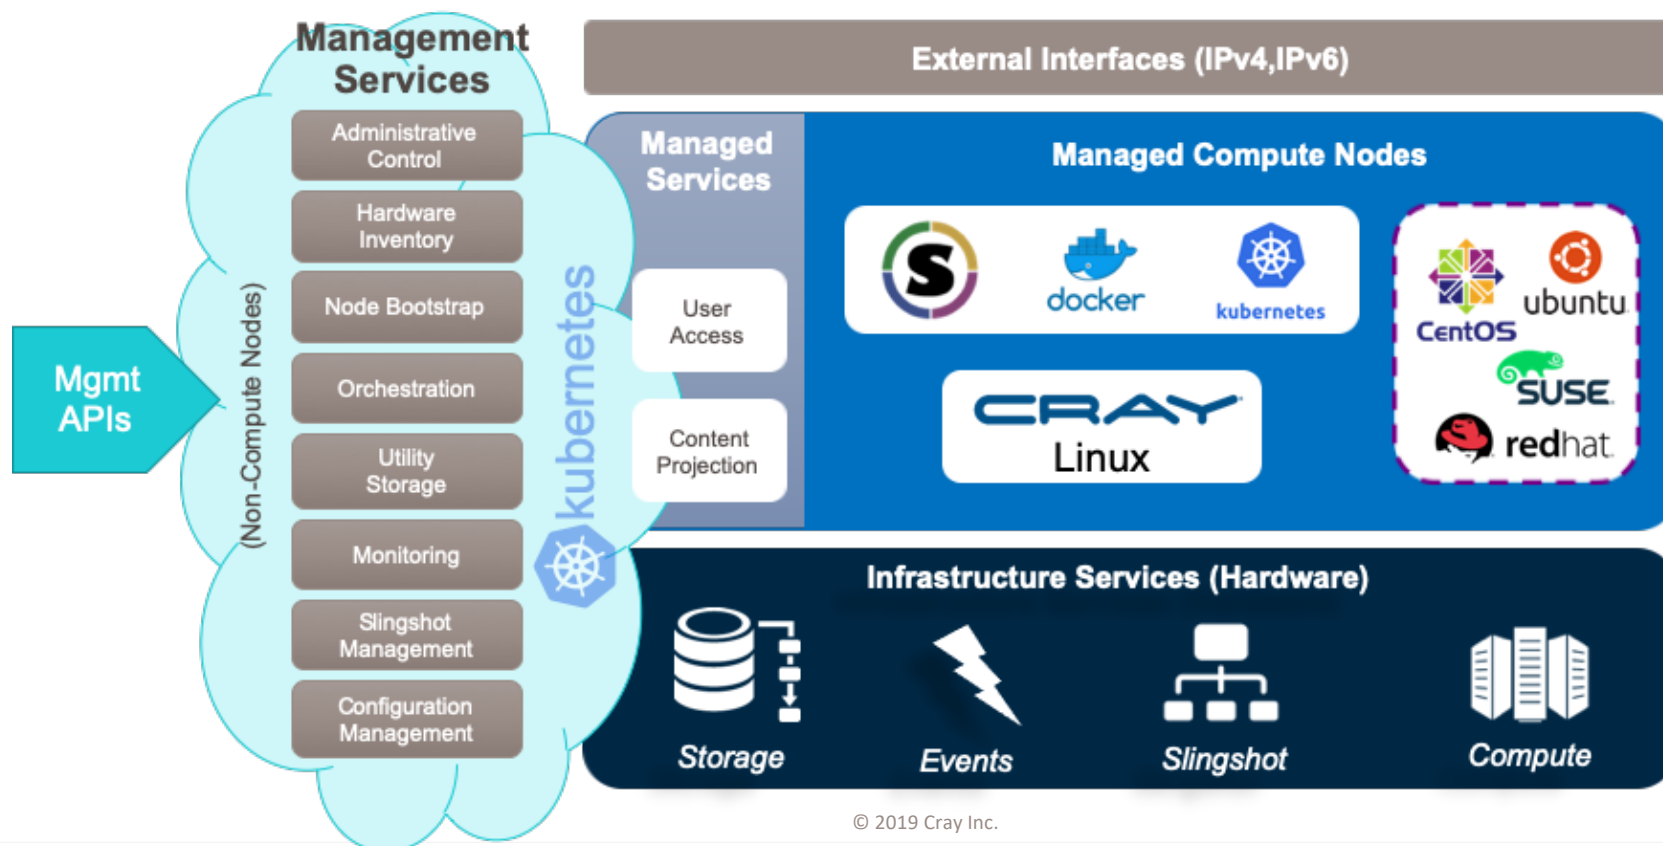


- El Capitan will meet its stockpile stewardship simulation mission
- System will feature:
 - Peak ≥ 2.0 DP exaflops
 - Peak power < 40 MW
 - Anticipating ~ 30 MW
 - Facility will support 85 MW total
 - Cray Slingshot interconnect
- Node will feature:
 - 4 to 1 GPU to CPU ratio
 - AMD Genoa CPUs based on Zen4
 - 3rd Gen AMD infinity architecture for high-bandwidth, low-latency connections

- Cray will provide several critical innovations
 - Cray and LLNL are working with ORNL jointly on non-recurring engineering (NRE) activities
 - Shasta software stack will feature greater flexibility under Cray's Compass program
 - El Capitan will include an innovative near node local storage solution

El Cap will on average provide 16 times more application capability on average vs. Sierra

El Capitan will employ Cray's Shasta system software



© 2019 Cray Inc.

El Capitan will exploit the flexibility of the Shasta Linux software stack to use RHEL



- Cray plans for fully optimized Linux for high-end HPC, based on SLES
 - Corresponds to current CLE software stack
- Individual Cray Software Components
 - Distro agnostic
 - Less intrusive, better interoperability with site software stack
 - Enables faster response time for updates
- Cray will provision for standard Linux distros with Cray network software
 - Possibilities include SLES, CentOS, Red Hat
 - Normal Cray image management and mapping tools can be used
 - Network stack is main component required – includes: NIC driver; Libfabric provider; & Lustre Network Driver
 - Other integration may optionally be supported
 - Heartbeat to management system
 - Simple REST calls via Node Management Network (NMN)
 - Blade power monitoring (requires Cray kernel module)
 - Access to system management APIs via NMN
 - Cray is also considering a middle ground with some Cray enhancements

LLNL will use Red Hat due its high quality QA processes and compatibility with TOSS

COE is a key part of our readiness activities and already in full swing.

- Porting to and testing the Cray SW stack
 - Largest piece of effort
 - Functionality and performance testing of compilers and tools
 - Open codes being put into a CI testing suite running on HPE machines to test very early SW
 - Restricted codes being tested at the labs
- Application porting to GPUs and HIP
 - Continue to move applications to GPUs
 - Work on RAJA and Kokkos HIP backends
- Application independent research for cross-cutting issues
 - GPU-aware MPI and messaging techniques
 - Strong scaling and impact of kernel launch
 - New hardware features coming in future GPUs





Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.